



ROD Newsletter

November 2011 Issue

The COD team: Malgorzata Krakowian, Marcin Radecki, Luuk Uljee, Ernst Pijper and Ron Trompert.

Introduction

This newsletter we have four topics lined up for you, namely, the Escalation Procedure, the Availability, Reliability and Unknowns, the NGI top-BDII availability and finally, the ROD metrics.

Input on the newsletter and Grid Oversight is very much welcomed by us. You can contact us by email at: central-operator-on-duty@mailman.egi.eu. A complete project glossary is provided at the following page: <http://www.egi.eu/results/glossary/>

Escalation Procedure

From time to time we received the signals from RODs that representation of escalation procedure in table form leaves too much space to own interpretation. To clarify the procedure we've prepared a flow diagram which is available here:

https://wiki.egi.eu/wiki/Operations:COD_Escalation_Procedure#Escalation_for_operational_problem_at_site

To see larger version of the graphics please click on the miniature.

If there is still something unclear with the procedure please let us know. We will try to answer on all your questions.

Availability, Reliability and Unknowns

Every month we, meaning the COD, are looking after the availability and reliability metrics of all sites. If sites fall below the 70%/75% availability/reliability threshold, we raise a GGUS ticket against those sites asking about the reasons for this. This is a very tedious and laborious task. To reduce the amount of work for us, for the sites and for yourselves we ask you be proactive and contact sites that are not doing too good in this field to help them solve their problems. Another thing is to encourage your site admins to look at your nagios a few times a day so site admins can detect solve problems before a ticket needs to be generated from the alarm. All of this contributes to provide users a better service, and of course, it makes your NGI look good.



We encourage you to have a look at: https://wiki.egi.eu/wiki/Availability_and_reliability_monthly_statistics. Here, under “Performance reports” you can find the league table where the performance of sites within your NGI is listed. This document is updated every month. For example, the league table of October is listed at: https://documents.egi.eu/public/RetrieveFile?docid=879&version=1&filename=EGI_League_WLCG_CREAM_LCGCE_CRITICAL_Oct2011.pdf

In the league table you see, in the section where the availability and reliability is given for each site within an NGI, a grey column listed as **UNKNOWN**. The number in that column gives the percentage of the month where the site was not monitored. The availability and reliability is computed over the period where a site is monitored. This means that when a site has an availability of 100% but the unknown percentage is 40%, then the site was 100% available over 60% of the month. This is not very meaningful because the actual availability could be anything between 60% and 100%. This means that NGIs should take care of the fact that the percentage of unknowns should be small.

Now there is a great variety in reasons why nagios reports unknowns. This could be due to problems at the sites but also problems in the NGI monitoring infrastructure. That’s why we are asking you to have a look at the league table for your NGI and have a look at the unknown percentage of your NGI. If the numbers are high for one site or on average over your whole NGI, it is worth taking a look into the cause of this problem.

Top-level BDII Availability

As you may know, the Resource Infrastructure Provider Operational Level Agreement has been approved by the Operational Management Board in October (<https://documents.egi.eu/document/463>). A new feature is the fact that a requirement of 99% availability has been imposed on the top-level BDII of NGIs.

You can have a look on these metrics of your own NGI at:

https://wiki.egi.eu/wiki/Availability_and_reliability_monthly_statistics#Performance_reports. .

We ask you or the people at your NGI responsible for the top BDII to have a look at https://documents.egi.eu/public/RetrieveFile?docid=879&version=1&filename=EGI-core_services_availability-per_NGI-Oct2011-1.pdf. If your NGI is in the yellow or red, think about on how to improve on the setup of your top BDII.

Currently, there are discussions going on on how to embed the monitoring of the NGI top BDII services into our monitoring infrastructure.



ROD metrics

This month for the first time based on data gathered on Operations Portal we send GGUS tickets to ROD teams which overlooked more than 10 issues on the dashboard.

It turned out that though we informed you about those metrics few times and though few means, there is still some problems with understanding the need of those metrics.

So:

1. ROD activity is one of the services which NGI is obligated to provide based on the agreement between NGI and EGI called RP OLA. It means that we need to be able to verify if the service is provided. ROD activity was established to ensure that operational problems are noticed and solved asap. Since the time is so important it was decided that escalation procedure which should follow ROD team will be strictly bounded with durations between following escalation steps. Based on this requirement ROD metrics was defined as a number of items which was not handled on time – alarms older than 72h and tickets expired.
2. Data for the metric are gathered by Operations Portal once a day at 4 AM and can be viewed here: <https://operations-portal.in2p3.fr/dashboard/metrics>
3. Item in the metric does not mean a particular not handled ticket or alarm but appearance of it in given day. It means that if you do not handle a ticket #1111 on Monday and on Tuesday it will be counted as 2 items. Why? Because it means that ROD team twice overlooked the issue and our goal is to count not site incidents but ROD incidents.
4. We set the target to 10 items because ROD activity is provided by people and we understand that human error may sneak into. We also understand national holiday's periods etc.
5. When you are above the target it does not mean explicitly that you are a bad ROD team. There may be reasons why the metrics is high. That is why we create the ticket to get this information.

NGI/ROC name	number of Tickets expired	number of Alarms > 72h	SUM
ROC AP	18	68	86
NGI_ARMGRID	10	65	75
NGI_UK	17	7	24
NGI_LA	22	0	22
NGI_DE	19	0	19
NGI_ME	15	2	17
NGI_CZ	0	14	14
NGI_IBERGRID	13	0	13
NGI_CANADA	9	0	9
NGI_FRANCE	1	8	9
NGI_IT	6	0	6
CERN	0	5	5
NGI_NDGF	5	0	5
NGI_NL	5	0	5
ROC RUSSIA	3	0	3
NGI_BG	0	2	2
NGI_IL	2	0	2
NGI_GRNET	0	2	2
NGI_SI	0	0	0
NGI_TR	0	0	0
NGI_RO	0	0	0
NGI_CH	0	0	0
NGI_PL	0	0	0
NGI_AEGIS	0	0	0
NGI_BA	0	0	0
NGI_BY	0	0	0
NGI_CYGRID	0	0	0
NGI_FI	0	0	0
NGI_GE	0	0	0
NGI_HR	0	0	0
NGI_HU	0	0	0
NGI_IE	0	0	0
NGI_MARGI	0	0	0
NGI_SK	0	0	0
NGI_IGLAC	0	0	0

Please note that the fact that the number of items on the COD dashboard is high for an NGI does not necessarily mean that the ROD team did a bad job. For example, this month it was noticed by NGI_UK that the dashboard counts also items during the weekends (thank you for noticed that!). We have created a ticket and ask to fix that. Next month if the problem will not be fixed, we will subtract such issues by hand so there might be some differences between what you get on Operations Portal page and in the newsletter table. Another example is for example NGI_IBERGRID. Here the amount of items on the COD dashboard was caused by a downtime of the NGIs nagios.